Localisation of a source of biochemical agent dispersion using binary measurements

Branko Ristic[a,1], Ajith Gunatilaka[a,*], Ralph Gailis[a,2]

[a] *Land Division, DST Group, Melbourne, Australia*

arXiv:1601.02304v1 [stat.AP] 11 Jan 2016

Abstract

Using the measurements collected at a number of known locations by a moving binary sensor, characterised by an unknown threshold, the problem is to estimate the parameters of a biochemical source, continuously releasing material into the atmosphere. The solution is formulated in the Bayesian framework using a dispersion model of Poisson distributed particle encounters in a turbulent flow. The method is implemented using the importance sampling technique and successfully validated with three experimental datasets under different wind conditions.

Keywords: Bayesian parameter estimation, binary sensor, turbulent dispersion, source localisation

1. Introduction

Localisation of a source of biological or chemical agent dispersing in the atmosphere is an important problem for national security and environmental monitoring applications [1]. Wind, as the dominant transport mechanism in the atmosphere, can generate strong turbulent motion, causing the released agent to disperse as a plume whose spread increases with the downwind distance [2]. Assuming a constant release of the contaminant, the problem involves estimation of source parameters: its location and intensity (release-rate). Two types of measurements are generally at disposal for source localisation: (i) the concentration measurements at spatially distributed sensor locations; (ii) the average wind speed and wind direction (typically available from a nearby meteorological station).

Many references are available on the topic of biochemical source localisation, assuming un-quantised (analog) concentration measurements. Standard solutions are based on optimisation techniques, such as the nonlinear least squares [3] or simulated annealing [4]. These methods are unreliable due to local minima or poor convergence; in addition, they provide only point estimates, without uncertainty intervals. The preferred alternative is the use of Bayesian techniques; they result in the posterior probability density function (PDF) of the source parameter vector, thereby providing an uncertainty measure to any point estimate derived from it. Most Bayesian methods for source estimation are based on Markov chain Monte Carlo (MCMC) technique, assuming either Gaussian or log-Gausiian likelihood function of measurements [5, 6, 7, 8]. Recently, a likelihood-free Bayesian method for source localisation was proposed in [9].

Binary sensor networks have become widespread in environmental monitoring applications because binary sensors generate as little as one bit of information. Such binary sensors allow inexpensive sensing with minimal communication requirements [10]. In the context of binary sensor networks, an excellent overview of non-Bayesian chemical source localisation techniques is presented in [11]. Best achievable accuracy of source localisation using binary sensors has been discussed in [12].

Prior work in using binary sensor data for biochemical source localisation assumes that the detection threshold of the sensor is known. It is a reasonable assumption for a commercial sensor whose sensitivity is specified (for example, in parts per million by volume (ppm_v) or grams per cubic meter) by the manufacturer and when the sensor is well calibrated. However, we consider at least two scenarios where the detection threshold of a binary sensor may not be accurately known. The first scenario is when a sensor's detection threshold goes off calibration due to environmental conditions such as temperature or humidity or ageing of the sensor. The second scenario is where the sensor is a human rather than a device. For example, imagine a person smelling a strong odour such as due to a gas leak or a decomposing animal carcass. When the person moves around, the smell will be detected in some locations but not in others, producing a binary measurement sequence without knowing the exact value of the threshold in ppm or g/m^3. In this paper, we develop a Bayesian algorithm that carries out source parameter estimation based on such *binary concentration measurements* where the sensor threshold is unknown. A Monte Carlo technique, importance sampling, is applied to calculate the posterior PDF approximately. The method is successfully validated using three experimental datasets obtained under different wind conditions.

*Corresponding author

Email address: ajith.gunatilaka@dsto.defence.gov.au (Ajith Gunatilaka)

[1] Now with RMIT University

[2] Now with NSID, DST Group

2. Models

2.1. Dispersion model

To solve the source localisation problem described above, we propose a solution formulated in the Bayesian framework which relies on two mathematical models: the atmospheric dispersion model and the concentration measurement model. A dispersion model mathematically describes the physical processes that govern the atmospheric dispersion of the released agent within the plume. The primary purpose of a dispersion model is to calculate the mean concentration of emitted material at a given sensor location. A plethora of dispersion models are in use today [13] to account for specific weather conditions, terrain, source height, etc. In this paper, we adopt a two-dimensional dispersion model of "particle encounters" in a turbulent flow, described in [14]. During a certain sensing period, each sensor experiences a Poisson distributed number of "encounters" with released particles. The binary nature of measurements indicates that a sensor reading of binary "1" or a "positive detection" corresponds to the number of such encounters exceeding a particular threshold.

If a binary sensor with a particular threshold makes positive detections (binary "1") at some locations and zero detections (binary "0") at other locations due to a source of a certain release rate, the measurements at these locations will be the same even if both the source release rate and the sensor detection threshold were scaled up or down together by the same amount; it is the ratio between the source release rate and the sensor threshold that determines which sensor locations will have positive or zero readings. Therefore, when we estimate the source parameters using binary data from a sensor whose detection threshold is unknown, it is not possible to estimate the absolute value of the source release rate; only the release rate normalised by the assumed sensor threshold can be estimated. Nevertheless, the source location, which is actually the parameter of main interest, can be estimated. Without loss of generality, in our experiments, we assumed the sensor to output binary "1" if it encounters at least one particle during a sensing period and output a binary "0" otherwise.

Let us assume that the biochemical source is located at (x_0, y_0), with a normalised release rate of Q_0. The particles released from the source propagate with the isotropic diffusivity D, but can also be advected by wind. We assume the released particles to have an average lifetime of τ. While the wind speed is typically available from meteorological data from a nearby measuring station, we use this speed as the prior guess for a Bayesian estimate of the true, effective wind speed affecting the advection of particles. Accordingly, let us assume that the mean wind speed is V and the mean wind direction coincides with the direction of the x axis. We denote the PDF of the wind speed by $\pi(V)$. A spherical sensor of small size a at a location with coordinates (x, y), non-coincidental with the source location (x_0, y_0), will experience a series of encounters with the released particles.

Denoting the parameter vector we wish to estimate, consisting of the source coordinates (x_0, y_0), normalised source release rate Q_0, and the wind speed V, by $\boldsymbol{\theta} = [x_0\ y_0\ Q_0\ V]^\intercal$, the rate of particle encounters by the sensor at the ith location (where $i = 1, \ldots, M$) with coordinates (x_i, y_i) can be modelled as [14]:

$$R(x_i, y_i|\boldsymbol{\theta}) = \frac{Q_0}{\ln\left(\frac{\lambda}{a}\right)} \exp\left[\frac{(x_0 - x_i)V}{2D}\right] \cdot K_0\left(\frac{d_i(\boldsymbol{\theta})}{\lambda}\right) \quad (1)$$

where D, τ and a are known environmental and sensor parameters,

$$d_i(\boldsymbol{\theta}) = \sqrt{(x_i - x_0)^2 + (y_i - y_0)^2} \quad (2)$$

is the distance from the source to ith sensor location, K_0 is the modified Bessel function of order zero, and

$$\lambda = \sqrt{\frac{D\tau}{1 + \frac{V^2\tau}{4D}}}. \quad (3)$$

2.2. Measurement model

The stochastic process of sensor encounters with released particles is modelled by a Poisson distribution. The probability that sensor at location (x_i, y_i) encounters $z \in \mathbb{Z}^+ \cup \{0\}$ particles (z is a non-negative integer) during a time interval t_0 is then:

$$\mathcal{P}(z; \mu_i) = \frac{(\mu_i)^z}{z!} e^{-\mu_i} \quad (4)$$

where $\mu_i = t_0 \cdot R(x_i, y_i|\boldsymbol{\theta})$ is the mean concentration at (x_i, y_i). Equation (4) represents the full specification of the likelihood function of parameter vector $\boldsymbol{\theta}$, given the sensor encounters z counts at the ith position.

However, because the actual sensor is binary, the measurement model is

$$b_i = \begin{cases} 1, & \text{if } z = 1, 2, 3, \ldots \\ 0, & \text{if } z = 0. \end{cases} \quad (5)$$

Note that b_i is a Bernoulli random variable with the parameter

$$\begin{aligned} q_i(\boldsymbol{\theta}) &= \Pr\{b_i = 1\} && (6)\\ &= \sum_{z=1}^{\infty} \mathcal{P}(z; \mu_i) && (7)\\ &= 1 - \mathcal{P}(0; \mu_i) && (8)\\ &= 1 - e^{-\mu_i}. && (9) \end{aligned}$$

The likelihood function for the sensor when it is at the ith location is then:

$$p(b_i|\boldsymbol{\theta}) = [q_i(\boldsymbol{\theta})]^{b_i} [1 - q_i(\boldsymbol{\theta})]^{1-b_i}. \quad (10)$$

Assuming sensor measurements are conditionally independent, the likelihood function of the parameter vector $\boldsymbol{\theta}$,

given the binary measurement vector $\mathbf{b} = [b_1, \ldots, b_M]^\intercal$, is a product:

$$p(\mathbf{b}|\boldsymbol{\theta}) = \prod_{i=1}^{M} [q_i(\boldsymbol{\theta})]^{b_i} \, [1 - q_i(\boldsymbol{\theta})]^{1-b_i}. \tag{11}$$

3. Parameter estimation

The estimation problem is formulated in the Bayesian framework. The goal is to compute the posterior PDF: the probability distribution of the parameter vector $\boldsymbol{\theta}$ conditional on the measurement vector $\mathbf{b}$. The posterior PDF provides a complete probabilistic description of the information contained in the measurements about the parameter vector $\boldsymbol{\theta}$. The basic elements required to compute the posterior distribution of are: (i) the prior distribution for the parameter vector $\pi(\boldsymbol{\theta})$ and (ii) the likelihood function $p(\mathbf{b}|\boldsymbol{\theta})$. Given these quantities, Bayes rule can be used to find the posterior PDF as

$$p(\boldsymbol{\theta}|\mathbf{b}) = \frac{p(\mathbf{b}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\int p(\mathbf{b}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}}. \tag{12}$$

The prior distribution $\pi(\boldsymbol{\theta})$ is typically non-Gaussian. For example, the source position can be restricted to polygon regions, while Q_0 and V are strictly positive random variables. Quantities of interest related to $\boldsymbol{\theta}$ (e.g., the posterior mean, variance) can be computed from the posterior PDF.

Optimal Bayesian estimation is generally impossible because the posterior PDF cannot be found in closed-form; this is certainly the case for the likelihood function specified in Sec.2 and a non-Gaussian prior $\pi(\boldsymbol{\theta})$. Hence we apply a Monte Carlo approximation technique of the optimal Bayesian estimation, known as importance sampling [15]. This technique approximates the posterior PDF by a weighted random sample $\{w_n, \boldsymbol{\theta}_n\}_{1 \le n \le N}$, which is created as follows. First, a sample $\{\boldsymbol{\theta}_n\}_{1 \le n \le N}$ is drawn from an importance distribution ϱ, i.e., $\boldsymbol{\theta}_n \sim \varrho(\boldsymbol{\theta})$, for $n = 1, \ldots, N$. The unnormalised weights are then computed as:

$$\tilde{w}_n = \frac{\pi(\boldsymbol{\theta}_n)}{\varrho(\boldsymbol{\theta}_n)} \, p(\mathbf{b}|\boldsymbol{\theta}_n) \tag{13}$$

for $n = 1, \ldots, N$. Finally, the weights are normalised, i.e., $w_n = \tilde{w}_n / \sum_{n=1}^{N} \tilde{w}_n$, for $n = 1, \ldots, N$. The choice of importance distribution ϱ plays a significant role in the convergence of point estimators based on the approximation $\{\boldsymbol{\theta}_n\}_{1 \le n \le N}$. Ideally ϱ should resemble the posterior. Since the posterior is unknown, good importance distributions are often designed iteratively (population Monte Carlo [15], progressive correction [16, 17]). As $N \to \infty$, however, the choice of ϱ is less relevant and even the prior π may be used as an (admittedly inefficient) importance distribution; this is done in our implementation for convenience. Furthermore, once the weighted random sample $\{w_n, \boldsymbol{\theta}_n\}_{1 \le n \le N}$ is computed, resampling (with replacement) is carried out [18] resulting in a sample with uniform weights. This last step was carried out mainly to improve the effect of visualisation of the posterior PDF (see figures in the next section). In our implementation, the prior PDF is adopted as:

$$\pi(\boldsymbol{\theta}) = \pi(x_0, y_0) \, \pi(Q_0) \, \pi(V) \tag{14}$$

where: $\pi(x_0, y_0)$ is a uniform distribution over designated polygon areas (e.g., buildings); $\pi(Q_0)$ is a Gamma distribution with shape k and scale parameter η; $\pi(V)$ is a normal distribution with mean $\bar{V}$ and standard deviation σ_V.

4. Experimental results

Three experimental datasets collected using a single moving binary sensor with unknown detection threshold were used in the paper to demonstrate the algorithm performance(further details of the experiments cannot be revealed on security grounds). In all cases, algorithm parameters were adopted as follows: $a = 0.2$ m, $D = 1$ m^2/s, $\tau = 1000$ s, $t_0 = 1$ s, sample size $N = 5000$, shape parameter $k = 3$, scale parameter $\eta = 7$, standard deviation of wind speed $\sigma_V = 0.2$ m/s. The wind conditions were different for the three datasets. The mean wind direction was specified by angle α, measured anticlockwise from the x axis. The values of wind parameters $(\bar{V}, \alpha)$ were $(0.28 \text{ m/s}, 195°)$, $(0.12\text{m/s}, -10°)$ and $(0.14\text{m/s}, 150°)$ for the first, second and the third dataset, respectively.

Fig.1 shows the results obtained using dataset 1: the left column displays the prior PDF $\pi(\boldsymbol{\theta})$; the right column - the posterior PDF $p(\boldsymbol{\theta}|\mathbf{b})$. Figs. 1.(a) and 1.(d) display the top down view of the area where the experiment was carried out. The placement and the readings of the binary sensor are also marked (the number of sensor locations in dataset 1 is $M = 45$). The locations where the sensor reported positive readings (i.e., $b_i = 1$), are marked by green squares; the remaining (non-detecting) sensor locations are marked by blue circles. Based on prior knowledge (e.g., intelligence reports), figures (a) and (d) also indicate the circular area where the source must be located (circle drawn in cyan colour). The center of this circle is the mean position of the sensor locations with positive readings; its radius is 150m. Furthermore, figures (a) and (d) also display the countours of the buildings (black lines). Both the marginal prior $\pi(x_0, y_0)$ and the marginal posterior PDF $p(x_0, y_0|\mathbf{b})$, are approximated by random samples marked by scattered red dots. Assuming the source must be inside one of the buildings, the marginal prior PDF $\pi(x_0, y_0)$ shown in figure (a) is a uniform PDF over the intersection of the cyan color bounded circular area and the area covered by the buildings. The marginal posterior PDF $p(x_0, y_0|\mathbf{b})$, shown in figure (d), concentrates in the upper left corner of one of the buildings, thus dramatically reducing the initial uncertainty in the source location. The true source location is marked by an asterisk at $(-298.4, -342.6)$. Figs.1. (b) and (c) show

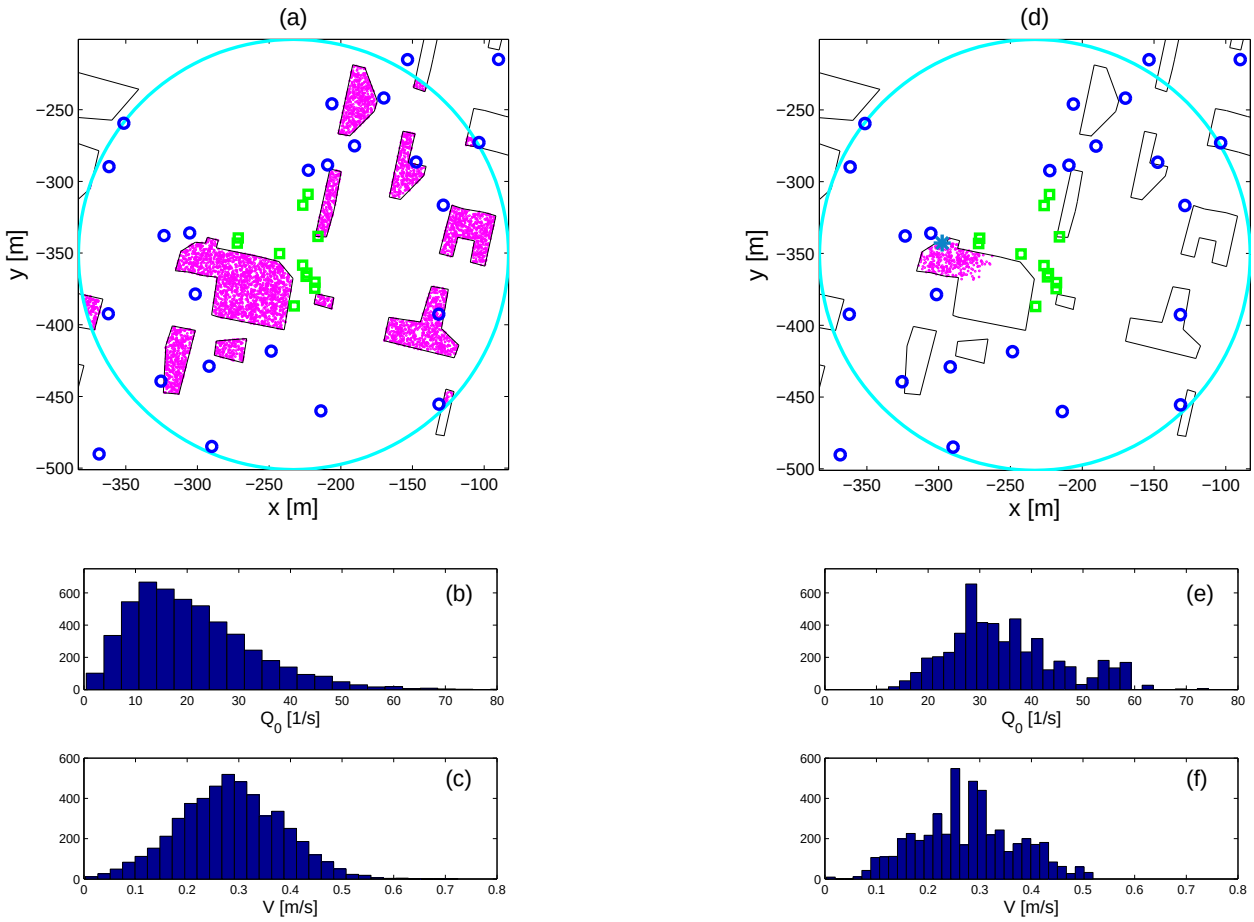

Figure 1: Estimation results for dataset 1: the left column is the prior PDF; the right column is the posterior PDF. Marginalised prior PDF $\pi(x_0, y_0)$ in (a) and the marginalised posterior PDF $p(x_0, y_0|\mathbf{b})$ in (d), approximated by random samples (indicated by scattered red dots). Sensor locations (green squares are positive readings, blue circles are non-detections), as well as the building contours (black lines), also indicated in (a) and (d). Wind direction coincides with the x axis. True source marked in (d) at $(-298.4, -342.6)$ (grey asterix). Figures (b) and (e) show the histograms corresponding to $\pi(Q_0)$ and $p(Q_0|\mathbf{b})$, respectively. Figures (c) and (f) show the histograms corresponding to $\pi(V)$ and $p(V|\mathbf{b})$, respectively.

the histograms of random samples approximating $\pi(Q_0)$; and $\pi(V)$, respectively. Figs. 1.(e) and 1.(f) display the histograms of random samples approximating the marginal posteriors $p(Q_0|\mathbf{b})$ and $\pi(V|\mathbf{b})$, respectively.

The marginal posterior PDFs, obtained by processing dataset 2, are shown in Fig.2. The number of sensor measurement locations in this case was $M = 25$. The initial random sample $\{\boldsymbol{\theta}_n\}_{1\leq n\leq N}$ was created in the same manner as for the case of dataset 1. From Fig.2.(a) we can observe that the posterior PDF $p(x_0, y_0|\mathbf{b})$ indicates fairly accurately the true source location marked by an asterisk at coordinates $(125.6, 436.6)$.

Finally, the marginal posterior PDFs, obtained by processing dataset 3, are shown in Fig.3. The number of sensor locations in this case was $M = 27$. The initial random sample $\{\boldsymbol{\theta}_n\}_{1\leq n\leq N}$ was created in the same manner as for the case of dataset 1. From Fig.3.(a) we can observe that the support of the marginal posterior PDF $p(x_0, y_0|\mathbf{b})$ indeed contains the true source location at coordinates $(31.2, -453.2)$. However, the conditions were such (the placement of sensor, the wind speed) that some ambiguity in the source location remains. The resulting posterior PDF is bi-modal (two buildings contain scattered red dots), suggesting that the source must be located in one of them.

5. Conclusions

The paper proposed a simple Bayesian estimation algorithm for localisation of a continuous source of biochemical agent dispersing in the atmosphere, using measurements collected at multiple locations by a single moving binary sensor whose detection threshold is unknown. The algorithm would also be applicable to a single snapshot of the measurements from a network of identical binary sensors. The sensor detection threshold may be unknown because it may have been drifted due to temperature, humidity, ageing, etc. Another possible scenario where the detection threshold of a binary sensor may be unknown is when a human, rather than a device, detects an odour at

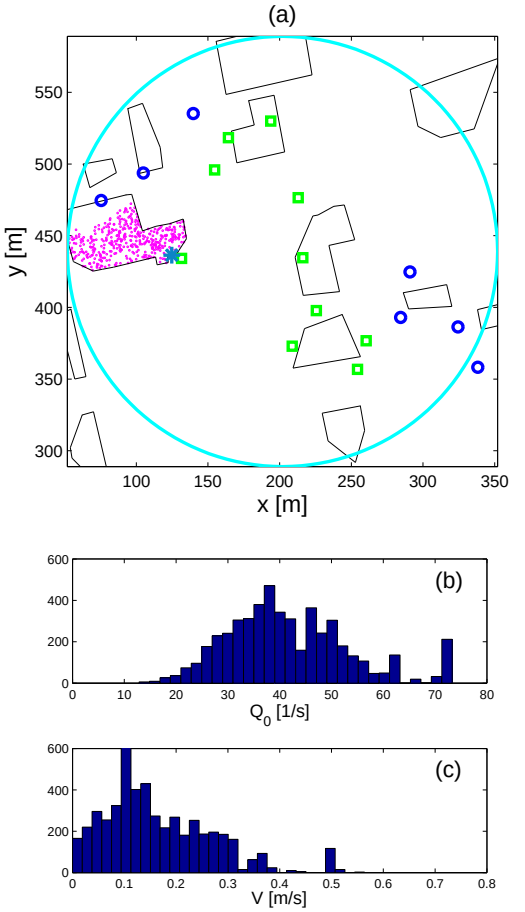

Figure 2: Estimation results obtained by processing dataset 2: (a) Marginalised posterior PDF $p(x_0, y_0|\mathbf{b})$ (scatter plot, red dots); sensor locations (green squares are positive readings, blue circles are non-detections); building contours (black lines); true source location at $(125.6, 436.6)$ (grey asterix). Wind direction coincides with the x axis. Figures (b) and (c) display the histograms of random samples approximating $p(Q_0|\mathbf{b})$ and $p(V|\mathbf{b})$, respectively.

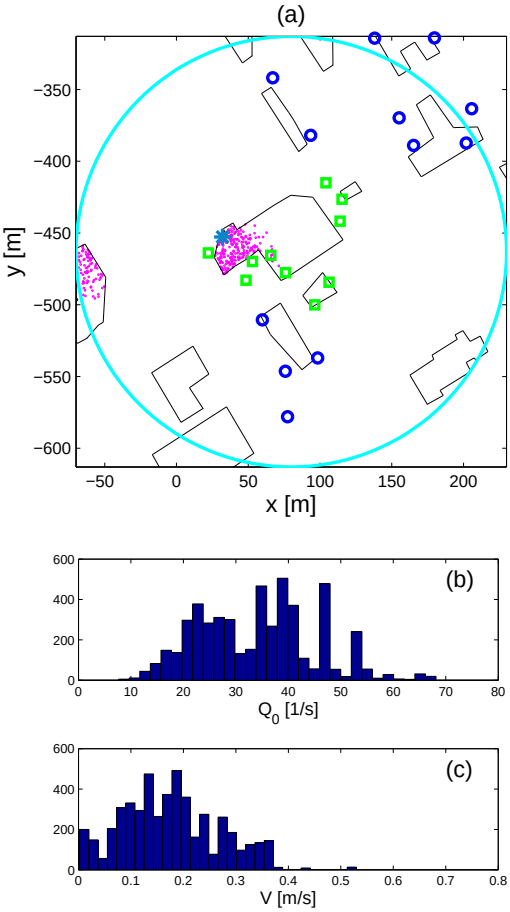

Figure 3: Estimation results obtained by processing dataset 3: (a) Marginalised posterior PDF $p(x_0, y_0|\mathbf{b})$ (scatter plot, red dots); sensor locations (green squares are positive readings, blue circles are non-detections); building contours (black lines). True source location at $(31.2, -453.2)$ (grey asterix). Wind direction coincides with the x axis.

some locations but not others. In this scenario, the person can easily make the binary measurements of "detection" or "non-detection" without knowing the exact detection threshold in terms of ppm or g/m^3 of the detected material. To enable source localisation in such scenarios, in our algorithm, we treat the source release rate, as well as the binary sensor threshold, as being unknown. Under these conditions, the algorithm can not estimate the absolute value of the source release rate and is only able to estimate the release rate normalised by the unknown sensor threshold. However, the algorithm can correctly estimate the location of the biochemical source. The performance of the algorithm is demonstrated using three experimental datasets collected in a semi-urban environment. In all three cases, the posterior density function included the true source location, thereby validating the proposed algorithm. Future work will consider introducing the uncertainty in the mean wind direction and more detailed dispersion models for urban environments.

References

References

[1] R. J. Kendall, S. M. Presley, G. P. Austin, P. N. Smith (Eds.), Advances in Biological and Chemical Terrorism Countermeasures, CRC Press, 2008.

[2] S. P. Arya, Pollution meteorology and dispersion, Oxford UniversityPress, 1998.

[3] J. Matthes, L. Gröll, H. B. Keller, Source localization by spatially distributed electronic noses for advection and diffusion, IEEE Trans. Signal Processing 53 (5) (2005) 1711–1719.

[4] L. C. Thomson, B. Hirst, G. Gibson, S. Gillespie, P. Jonathan, K. D. Skeldon, M. J. Padgett, An improved algorithm for locating a gas source using inverse methods, Atmospheric environment 41 (2007) 1128–1134.

[5] A. Keats, E. Yee, F.-S. Lien, Bayesian inference for source determination with applications to a complex urban environment, Atmospheric Environment 41 (3) (2007) 465 – 479.

[6] R. Humphries, C. Jenkins, R. Leuning, S. Zegelin, D. Griffith, Atmospheric tomography: a Bayesian inversion technique for determining the rate and location of fugitive emissions, Environmental Science and Technology 46 (3) (2012) 1739–1746.

[7] M. Ortner, A. Nehorai, A. Jeremic, Biochemical transport modeling and Bayesian source estimation in realistic environments, IEEE Trans. Signal Processing 55 (6) (2007) 2520–2532.

[8] I. Senocak, N. W. Hengartner, M. B. Short, W. B. Daniel, Stochastic event reconstruction of atmospheric contaminant dispersion using Bayesian inference, Atmospheric Environment 42 (33) (2008) 7718 – 7727.

[9] B. Ristic, A. Gunatilaka, R. Gailis, A. Skvortsov, Bayesian likelihood free localisation of a biochemical source using multiple dispersion models, Signal Processing 108 (2015) 13–24.

[10] J. Aslam, Z. Butler, F. Constantin, V. Crespi, G. Cybenko, D. Rus, Tracking a moving object with a binary sensor network, in: Proc. 1st Int. Conf. on Embedded Networked Sensor Systems, SenSys '03, ACM, New York, NY, USA, 2003, pp. 150–161.

[11] H. Chen, Greedy methods in plume detection, localization and tracking.

[12] B. Ristic, A. Gunatilaka, R. Gailis, Achievable accuracy in gaussian plume parameter estimation using a network of binary sensors, Information Fusion 25 (2015) 42–48.

[13] N. Holmes, L. Morawska, A review of dispersion modelling and its application to the dispersion of particles, Atmospheric Environment 40 (2006) 5902–5928,.

[14] M. Vergassola, E. Villermaux, B. I. Shraiman, 'Infotaxis' as a strategy for searching without gradients, Nature 445 (25) (2007) 406–409.

[15] C. P. Robert, G. Casella, Monte Carlo Statistical Methods, Springer, 2004.

[16] C. Musso, N. Oudjane, F. LeGland, Improving regularised particle filters, in: A. Doucet, N. deFreitas, N. J. Gordon (Eds.), Sequential Monte Carlo methods in Practice, Springer-Verlag, New York, 2001, Ch. 12.

[17] M. Morelande, B. Ristic, Radiological source detection and localisation using bayesian techniques, IEEE Trans Signal Processing 57 (11) (2009) 4220–4231.

[18] B. Ristic, S. Arulampalam, N. Gordon, Beyond the Kalman filter: Particle filters for tracking applications, Artech House, 2004.